

Getting Algorithm Visualizations into the Classroom

Clifford A. Shaffer
Dept. of Computer Science
Virginia Tech
shaffer@cs.vt.edu

Monika Akbar
Dept. of Computer Science
Virginia Tech
amonika@vt.edu

Alexander Joel D. Alon
Dept. of Computer Science
Virginia Tech
alon@vt.edu

Michael Stewart
Dept. of Computer Science
Virginia Tech
cse@vt.edu

Stephen H. Edwards
Dept. of Computer Science
Virginia Tech
edwards@cs.vt.edu

ABSTRACT

Algorithm visualizations (AVs) are widely viewed as having the potential for improving computer science education. However, the rate of AV use and overall impact on education does not match the positive interest in their use that instructors report. Surveys of CS faculty show that impediments to successful use of AVs in the classroom include difficulties in finding quality AVs on desired topics, difficulties in adapting AVs to a given classroom setting, and lack of knowledge on the best way to deploy AVs. This indicates a need for better support for instructors, to get them past these barriers. We seek to provide this support through an online educational community that relies on a new model based less on the “digital library” approach of information gained by going to a site and searching. Instead, the focus is on community-added content through members’ discussions, reviews, and ratings of content items. The AlgoViz community effort will better focus the future direction of AV development and use.

Categories and Subject Descriptors

E.1 [Data Structures]; K.3.2 [Computers and Education]: Computer and Information Science Education

General Terms

Algorithms

Keywords

Algorithm visualization, data structure visualization, online educational community, active learning, animation, interactive tutorial, teaching aid

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCSE'11, March 9–12, 2011, Dallas, Texas, USA.

Copyright 2011 ACM 978-1-4503-0500-6/11/03 ...\$10.00.

1. INTRODUCTION

Algorithm (and data structure) visualizations (AVs) are widely viewed as having the potential for improving computer science education. Surveys of instructors over the past decade have shown great support for the concept of AV use in the classroom. Yet the rate of AV use and overall impact on education does not match the interest claimed by instructors. These same surveys show that the impediments to successful use of AVs are related to difficulties in first finding good materials, and then difficulties in deploying them in classes effectively. Fortunately, much progress has been made in recent years. There are many high-quality AVs available—though there are enough poor quality AVs that sorting out the good ones can be time consuming [16]. There is a growing body of research on how to use AVs effectively in the classroom (though getting up to speed on this information is again time consuming). These impediments could be overcome with adequate information about which AVs are of good quality and how they can best be used. From the developers’ side, more information about what makes for an effective AV, and which topics are of greatest need, could lead to more effective AVs being produced.

Thus, we find that information and dissemination of best practice plays a key role in the successful deployment of AVs in the classroom. Given that the most needed information is (1) feedback on AV quality (user ratings) and (2) hearing of real experiences with this material in the field, the most effective way to improve instruction is through an active community of AV users and developers interacting to exchange information and experiences.

This paper is organized as follows. Section 2 discusses past surveys of instructor opinions on AVs, and presents the results of a new survey given during SIGCSE’10. Section 3 discusses the role that an online educational community can play in overcoming barriers to adoption of AVs. Section 4 describes the AlgoViz Portal. Section 5 describes our efforts related to moving from a digital library model to an online educational community model. Future plans are discussed in Section 6. Readers interested in background related to the history AVs, their level of availability, and best practices for their use in the classroom are referred to [16, 9, 7].

2. SURVEY RESULTS: A DISCONNECT

Periodic surveys on both instructor interest in AVs and their level of use have been conducted among attendees of education conferences and listservs for more than a decade.

Naps, et al. [9] report on three surveys conducted in 2000. Collectively, they indicate a strong positive view in favor of AV use (over 90%). However, only about 10% of respondents indicated frequent use of AVs in data structures courses at their institutions, while about half to three quarters indicate occasional use of AVs in such classes.

At SIGCSE'10 (held in Milwaukee during March 2010) we conducted a new survey to determine instructor attitudes toward AVs and their use in the classroom. The survey was designed in the spirit of the 2000 surveys reported in [9]. The survey instrument was included in the conference bag distributed to SIGCSE'10 attendees, it was distributed at the AlgoViz Project booth in the exhibition hall, and it was also handed out and collected at a special session [17] on the AlgoViz community effort.

Survey participants (both of our survey and the previous surveys reported) do not reflect a random sample of CS educators. First, they were attendees of various CS education conferences. Second, they probably had some interest in algorithm visualization, since they voluntarily returned the survey. In our case, often this was triggered by a visit to our booth or the special session. Others might have been attracted by the opportunity to win a \$25 gift card to Amazon.com that was advertised as a mechanism to increase response rates. Yet, the survey does provide a snapshot of the issues that concern CS instructors regarding AV use, and it serves as a useful comparison to the previous surveys (of similar populations) to see whether anything has changed.

Question 1: Using algorithm visualizations can help learners learn computing science.

- Strongly Agree: 18
- Agree: 23
- Neutral: 2
- Disagree: 0
- Strongly Disagree: 0

This question closely mirrors questions asked in the prior surveys. The results are pretty much the same as 10 years ago. In these and all other interactions that we have had with instructors over the years, they have consistently and overwhelmingly indicated that they think using algorithm visualization in courses is a good idea.

Question 2: Have you used any AVs in a class that you taught in the past two years?

- Yes: 22
- No: 19

At a coarse level, this is not much different from results reported in 2000. The 2000 survey was finer-grained in that it asked for information about level of AV use. The 2000 survey reported about 10% of instructors using AVs in relevant classes on a weekly basis, and over 50% using AVs on at least an infrequent basis (but there is no time period defining “infrequent”). It is not clear that all of the “yes” answers from the present survey refer to what we might consider to be “AVs.” But they probably all refer to some sort of dynamic software visualization run on a computer, as was the case in the 2000 survey.

Important technological factors contribute to making AVs easier to use in the classroom than in prior years. More AVs are available than ever before, backed up by improved research studies and improved access both through general Internet search and at the AlgoViz Portal (see Section 4).

Increased access to the Internet by students and instructors, both in and out of the classroom, makes AV use more practical. For example, at Virginia Tech, while it has been “possible” to project Internet material from a computer in a class for over a decade, it has only been in the past couple of years that such access has become both ubiquitous and reliable in all of our classrooms. This makes a huge difference in the confidence of mainstream instructors for using such technology, in contrast to early adoptors [4]. From discussions with instructors and from comments on the survey forms, we know that even now not all schools have built up the proper classroom support infrastructure to allow broadcasting of internet presentations in typical classrooms. This is still an evolving technology.

Question 3: What do you see as the greatest impediments to using AVs in your courses?

We have aggregated the text responses into the following broad categories.

- Trouble with finding suitable AVs: 13
- Trouble with integrating AV material into course: 11
- Difficulty of making AVs: 3
- Finding class time to fit AVs in: 2
- Classroom technology limitation: 1
- Question the pedagogical utility of AVs: 1

These results are roughly the same as reported in [7]. We see here that the largest single category is trouble with finding materials. Interestingly, the need for an organized repository for AVs (with some review mechanism) has been recognized since at least 1996 [1]. We can hope that the presence of newer online resources such as the AlgoViz Portal will improve this situation. The second largest category is much harder to deal with. It relates to integrating materials into courses once they have been discovered or created. This is a well-known problem [4]. While it is easy to provide pointers to AVs to students as supplemental material, it is much harder to integrate AVs into lectures, labs, and assignments.

3. AN ONLINE COMMUNITY

As indicated in the surveys, the greatest impediment beyond finding suitable AVs is integrating them into courses. If we hope to increase the penetration of AVs into the classroom, then instructors will need sufficient support to get over this hurdle. While improved access to information about what AVs are available can help, something beyond this seems needed to complete the circle. Instructors need information about how to use AVs in class, and confidence that their use will lead to success. This sort of information will need to come from other instructors as well as education researchers.

The central idea of the AlgoViz Portal is to prototype a new use model for online educational sites, with the goal of establishing an online educational community centered around Algorithm Visualization. The Portal diverges from the old model of an information site as a “digital library” where people come to search a collection. The problem with the traditional model is that users do not get the interaction with other instructors that they need to build confidence in the artifacts. Instructors do not get necessary information about which artifacts are truly of value, and how best to use them. Another problem is that users do not know when new content of interest has been added to the site. One must also

recognize that it is difficult for content providers today to improve over the effectiveness of modern search engines for discovering online content. Static digital library “services” typically can be provided by search engines automatically.

The new model embraced by the AlgoViz Portal is to allow the community to add value beyond the base content. This comes as a byproduct of direct communication between community members. While forums are one obvious method for this, for many communities there is a deeper communication involving evaluation of the content (community ratings and recommendations for content entries), and discussions about the content (sharing experiences on how to make the best use of it). This user-based, value-added material is not “content” under the old model, and is generally not captured by the metadata associated with a collection. Yet, in many educational contexts, the user feedback on items in the collection is as important as the collection itself. The reason is that in educational communities, a major concern for instructors is deciding whether a given educational resource is of good quality, and how to use it.

Our analogy is to Amazon.com. While its backbone is a catalog and purchasing infrastructure, Amazon is much more than this from a community standpoint. An important component to Amazon’s success is community-generated ratings and reviews, where significant categories such as books, movies, and music have increased sales by as much as 20% due to the value added by these reviews [18]. In the context of educational content, a successful online educational community is much more than a digital library. It must also include community-added content that raises the value of the information. This includes discussion of best practice, and user evaluation for the artifacts represented in the library.

Due to the size of its user base, Amazon has a rich collection of user-supplied product ratings and reviews even though less than one tenth of one percent of their customers provide this information [18]. An online educational community for a topic such as AV use has an additional burden, since it must get a much higher percentage of the potential “market” to actively engage, more like 1 to 10% of the membership, depending on how membership is defined. Fortunately, members of that “market” also have a greater incentive to become active participants.

4. ALGOVIZ.ORG

The *AlgoViz Portal* (<http://algoviz.org>) is intended to be an online hub for educators, developers, and users of algorithm visualizations. It is the successor of the AlgoViz Wiki [20], an older project by the Virginia Tech Algorithm Visualization Research Group.

AlgoViz.org is built on Drupal [2], an open-source content management system written in PHP. The functionality built into the basic Drupal installation, known as Drupal core, can enable a website with features such as forums, blogs, news feeds, comments, and user registration. Perhaps the most important consideration in the selection of Drupal for the CMS powering the AlgoViz Portal is the massive existing body of user-contributed modules, which greatly reduces development time, combined with the ability to create new modules or modify the existing code base if necessary. AlgoViz.org is built on top of a Drupal distribution, Acquia, that packages Drupal core with many popular contributed modules. For its database backend, AlgoViz.org uses a MySQL

Virginia Tech - Interactive Hashing Tutorial

View Edit Outline Track Visitors

Link(s) <http://research.cs.vt.edu/AVresearch/hashing>

Topic(s) Hashing, Search Algorithms

Recommendation Recommended

Works? Yes

Delivery Method(s) Java Applet

Project Virginia Tech Algorithm Visualizations

Project Relationship Part of collection

Language(s) English

Author(s) Mayank Agarwal, Matt Jeswa, Arpit Kumar, Purvi Saraiya, Crystal Weil, A.J. Alon, Cliff Shaffer

Institution(s) Virginia Tech

Activity Level(s) Exploration, Questions, Random data, Step control, User data

AV is good for Comparison, Lab exercise, Lecture aid, Self-study

Source Code License Licensed under GPL

Last Modified N/A

DESCRIPTION

A complete tutorial for teaching hashing, largely based on the presentation in "A Practical Introduction to Data Structures and Algorithm Analysis" by Clifford A. Shaffer. A series of HTML pages takes the reader through the topics of hash functions, open vs. closed hashing, collision resolution, deletion. The web pages are interspersed with visualizations that demonstrate a variety of hash functions and collision resolution methods, as well as applets that show aspects of the performance for the various methods.

EVALUATION

The presentation is fairly detailed. This tutorial can stand alone as an introduction to hashing, and has been classroom tested in that mode. The series of applets that show the various hash functions and collision resolution methods allow students (optionally) to working in "test" mode where they must first predict where the next key will go in the hash table. The applets that demonstrate performance are

Figure 1: A sample catalog entry.

database driven by the PHP MySQLi (MySQL Improved) Extension.

The main informational resource at the Portal is the *AlgoViz.org Algorithm Visualization Catalog*, a referatory to AVs on the Web. It is not intended to be merely a simple link collection. Each catalog entry contains key details about the AV such as its delivery mechanism, information on the developers and the project (if any) that the AV is part of, an objective description, and a subjective evaluation of the pedagogical value of the AV (see Figure 1). The AV Catalog has expanded in size and scope since its original incarnation as part of the AlgoViz Wiki, and now contains over 500 entries.

Figure 1 shows a sample entry in the Catalog. Additional content includes a description of Activity Level, which specifies how a user may control a particular visualization. The References field is used to list links to academic publications or white papers associated with a particular AV. Users are encouraged to provide Digital Object Identifier (DOI) links as references whenever available. Entries may also include multimedia in several forms: attached Screenshots, embedded Videos, or links to external galleries. Such material would ideally demonstrate how an AV works or even classroom use cases.

The other major informational resource at the Portal is an *annotated bibliography* of the published literature on algorithm visualization and related research. An entry in the bibliography contains information about that publication, including ISSN, ISBN and DOI numbers when available. This information is presented to the user in a tabular format, but may also be displayed in Tagged, XML, or BibTeX form. A link is provided to Google Scholar for most entries. The annotated bibliography currently contains over 500 papers.

The Portal provides several public forums for discussion of AV-related topics. We hope that the forums will be the primary gathering place for members of the AlgoViz.org community, particularly those who are new to the field or who have specific questions.

Instructors report that a major deterrent to their use of AVs is the time and effort they perceive must be expended in integrating existing AVs into existing course content (see [9] and the survey results in Section 2). To help address this

issue, the Portal provides a structure called the *field report*, which gives a convenient way for instructors who have used an algorithm visualization in a class to share their experiences. Field reports are designed to supplement the evaluation data in the AV Catalog, since AV ratings in the catalog are not necessarily based on real-world classroom experiences. A field report can thus provide empirical evidence to strengthen or qualify a recommendation. Instructors submitting a field report are asked to describe the lesson in question and how AVs were incorporated into the lesson. They may assess how successful the lesson was as they see fit. They are also asked to provide the following:

- the person or persons who presented the lesson (in case the user submitting the field report is not the same as the presenter(s));
- the institution where the lesson took place; and
- the AV(s) used (as a link in the catalog).

The AlgoViz Project is a primary content contributor to Ensemble [3], the NSDL Computer Science Digital Library. To ensure interoperability between the different repositories in its index, the Ensemble project employs the Open Archives Initiative Protocol for Metadata Harvesting, or OAI-PMH [11]. Under OAI-PMH, actors fall into one of two roles, the data provider and the service provider. A data provider is a repository or collection that structures and exposes contained metadata. A service provider requests metadata via service requests sent to data providers. We have implemented for the AV Catalog an OAI-PMH data provider service that maps catalog entry content to tags in the Dublin Core format. The metadata is thereby made available to Ensemble and other potential OAI service providers.

AV developers can benefit from a source of example AVs to use for reference, including source code. We have established OpenAlgoViz at SourceForge.net [12]. The site has at least two purposes: to provide reference implementations for good quality AVs that others can use to see how things should be done, and to provide a hosting site for AV contributors. OpenAlgoViz has the potential to become a major repository of AV source code in its own right. All entries are open source and are contributed by developers in the AV community. Currently, the repository contains the following contributions: JFLAP from Susan Rodger of Duke University [14]; several algorithm visualizations from Virginia Tech; the JHAVÉ [8] and GAIGS [10] projects from Tom Naps of University of Wisconsin–Oshkosh, the PCIL project [6] from Mississippi State and Tennessee Tech, internationalization libraries [15] from Guido Rößling of Technische Universität Darmstadt, and XAAL, an XML-based markup language for algorithm animation [5], from Ville Karavirta of Aalto University in Finland.

We routinely monitor AlgoViz Portal site traffic to view user trends using both Google Analytics and native Drupal tools. Since the portal’s inception in July 2009, we have received more than 9,800 visits from outside the Virginia Tech domain. Figure 2 shows the monthly visits from July, 2009 to August, 2010. Monthly traffic at the portal has increased by nearly a factor of four during this time. We observed a major spike in activity caused by advertising the site during the SIGCSE 2010 conference. Unfortunately, site traffic data must be interpreted in light of the fact that our site has been split between the AlgoViz Portal and the original AlgoViz Wiki. Until June, 2010, the Wiki contained the

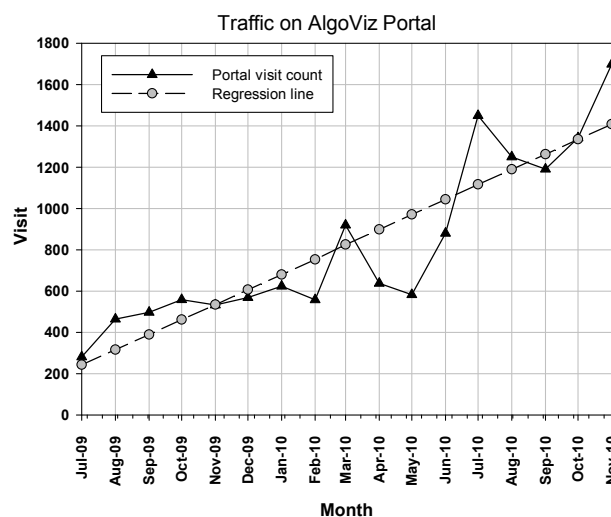


Figure 2: Visitors per month to the AlgoViz Portal, July 2009 (shortly after the Portal was made public) to November 2010. A regression line shows the overall trend.




Figure 3: Monthly page views and bounce rates of the AlgoViz Portal, July 2009 to November 2010.

actual catalog (presumably a primary attraction for users). Any visitor moving from the Portal to the Wiki would be considered as a visitor who had left the Portal. A sustained increase in traffic measured for the Portal began in June 2010, when the Wiki catalog was migrated to the Portal.

Figure 3 shows traffic data in terms of page views (the loading of a new page). The average number of page views per visit is about three. The total number of page views is steadily increasing over time. The bounce rate is an estimate of how many visitors left the site after viewing a single page; in most cases that could be the front page. The bounce rate now averages about 54%.

Since July 2009 we have had over 9,000 unique visitors to the Portal. 64% of the visits came from first-time visitors. This could be due to the relative newness of the site. Because of our activities in conferences and public mailing lists, we are likely to get more new visitors at this point. On average, new visitors spend a little more than three minutes in the site while returning visitors spend six minutes.

5. BUILDING THE COMMUNITY

Since a relatively high percentage of the community must actively engage to reach the needed critical mass to sustain community interaction, much of our work focuses on ways to lower the barriers to participation. This includes technology to support easy/automatic login from allied sites (to avoid additional registration or login overhead), a robust notification system to inform users when new information of interest has been provided at the site, and mechanisms to solicit ratings information.

5.1 Community Value Added

An important example of community-added content is a community rating system for AVs in the catalog. Currently, quality assessment information in the catalog comes from:

- An overall ranking given by the catalog editors (Recommended, Has Potential, Not Recommended) on each of four potential uses: lecture aide, self-study supplement, stand-alone use, and debugging aide;
- ratings by users on a five-star basis, similar to product ratings at a site such as Amazon;
- description and review information for the AV;
- field reports; and
- user comments.

While the present incarnation of the catalog gets its primary review information from the catalog editors, in the future the bulk of this material must come from the community to reach its true potential, through users submitting reviews and comments. This in turn will drive recognition of AVs and other content items that can then be sent as “notifications” to community members through their favored notification streams, which will in turn prompt community members to add more value to the community through their own ratings and comments.

5.2 AlgoViz Awards

In 2010 we initiated the AlgoViz Awards, an annual competition intended to provide a number of positive effects. Our initial goal was to help build community interest in the site, and to showcase high-quality AVs. Each year, the winners of the annual awards will be added to our AV Hall of Fame. This should serve as a useful starting point for instructors who want to see what is possible with AVs.

The AlgoViz Awards began with a nomination process starting in June 2009 and closing in October 2009. Any AlgoViz Portal registrant was permitted to nominate an entry. Several AVs were selected by the AlgoViz Project Steering Committee to initially pre-seed the competition. Nominees were assigned to one of three categories: *Standalone AVs in Lower Division Topics* represents AVs that cover topics generally taught in CS1, CS2, and sophomore-level courses; *Standalone AVs in Advanced Topics* represents AVs that cover topics taught in upper division and graduate-level courses; *Online Tutorials and Hypertextbooks* represents AVs that are integrated into a larger learning object, instead of presented alone. These award categories were chosen in part to draw attention to various issues in AV development. In total, fourteen AVs were nominated for an AlgoViz.org Award.

All registrants at the Portal were given the opportunity to vote on the nominated visualizations. The actual voting process involved rating each AV on a five-point scale, where

one point represents a poor AV and five represent an excellent AV. Users were permitted to rate as many or as few AVs as desired, and were allowed to change their rating choice on any AV as many times as desired until voting closed (only the final rating on each AV was counted). Voting ended on 31 January 2010. The winners were publicly declared in a video presentation on 13 March 2010 at the SIGCSE conference. The video gave a quick demonstration showcasing the strengths of each winning AV. It is available at http://www.youtube.com/watch?v=QLT_AKgYJYk. The second annual AlgoViz Awards includes a category on program visualization tools, and a category on new AVs. Results will be announced during SIGCSE’11.

5.3 Notification

An important way to keep members in touch with the site—and thereby increase participation—is through notification mechanisms. We cannot expect users to routinely come to our site to determine if new content is available, since the overhead for doing so is too high relative to the frequency of change. Instead, we need to notify members of our community whenever “interesting” content appears.

This notification needs to be done in a way that is natural to any given individual, in that it fits in with their normal information flow. For some, that flow takes the form of email or RSS feeds. For others, notification streams come through social networking schemes such as Twitter or Facebook. Each user must have sufficient control in determining what they consider to be “interesting” so that they do not receive too much unwanted information.

Many people opt to use social networking schemes for notification, rather than receive all such information via more traditional email, because those social networking systems are better tuned toward integrating many streams of small, targeted notification (a very short message, typically). As a consequence, notification streams based on changes to the AlgoViz site content need to be fine grained, and targeted to only those recipients that want any given notification. For example, a given user might want to know when certain AV topics are updated in the catalog, or new posts are made to certain subforums at the site, and not others. Thus, we are building an extensive “interest profile” mechanism whereby users can easily select what topics are of interest, and the method by which they wish to be notified about new content. For now, users can control their level of notification by selecting from various twitter or RSS feeds.

6. FUTURE WORK

People are more willing to visit online communities where there is more activity, and are more willing to contribute to communities where they see contribution is the norm [13, 19]. This would seem to argue for minimizing barriers to contribution. Forcing people to register initially, and sign in on return visits, can be a significant barrier to getting started for many people. Many of us are reluctant to invest the overhead of time to register, or give any personal information (especially an email address, which is often the minimum price of admission). So why require this?

Unfortunately, there are many malicious entities on the Internet. The most common is the “link spammer” who attempts to hijack the public’s ability to deliver content to a website in order to promote links to unrelated sites, with the goal of boosting the 3rd party site’s visibility in search en-

gines. If permitted, these sorts of attacks can dominate the “contributed” content at the site and destroy its usefulness to the targeted community.

For this reason, AlgoViz currently requires potential contributors to register first. But since encouraging content contributions (such as ratings, reviews, and forum posts) is of primary concern, and since registration is a major disincentive to “simple” contributions such as rating an AV, we see an urgent need for lowering this barrier. One way to minimize the need for registration is to accept authentication from other sites. The key information that we actually need is some unique identifier for each visitor, so that we can recognize them as the same person at each return visit. In that way we are able to connect the contributions of the individual over time, and to block malicious individuals. Thus, explicit registration at a community site is not absolutely necessary to its goals. An alternative is shared registration with explicit allied sites (such as the Ensemble portal [3]), OpenID access (not as useful since it still requires initial registration), and automated connection from sites such as Facebook or LinkedIn.

Another approach used by some sites is relaxed registration combined with greater moderation. That is, we could eliminate the requirement for validation by returning an email message. But unvalidated users are the greatest risk for malicious behavior. By sending all contributions by unvalidated users to a moderator before making the contribution public, we can screen these attacks out. Once a user has cleared some number of successful moderation attempts, they can then be flagged as a “trusted” user with no further need for automatic moderation.

We will also study the potential for “social capital” to encourage users to contribute. Many sites are successful by giving some sort of “reward” for contribution in the form of titles in a hierarchy. Those who contribute more get better titles in recognition. While this is seemingly minor, it appears to tap into the innate competitiveness by some people, and the need for validation of authority by others.

7. ACKNOWLEDGMENTS

We gratefully acknowledge support of the AlgoViz Project by the National Science Foundation’s CCLI and NSDL programs under grants DUE-0836940 and DUE-0937863.

8. REFERENCES

- [1] J. Bergin, K. Brodie, M. Patiño-Martínez, and seven additional authors. An overview of visualization: its use and design: Report of the working group in visualization. In *ITiCSE '96: Proceedings of the 1st Conference on Integrating Technology into Computer Science Education*, pages 192–200, 1996.
- [2] Drupal. <http://drupal.org>, 2010.
- [3] ENSEMBLE. <http://www.computingportal.org/>.
- [4] K. Hew and T. Brush. Integrating technology into k12 teaching and learning: current knowledge gaps and recommendations for future research. *Educational Technology Research and Development*, 55:223–252, 2007.
- [5] V. Karavirta. Xaal - extensible algorithm animation language. Master’s thesis, Helsinki University of Technology, Helsinki, Finland, Dec 2005.
- [6] B. Malone, T. Atkison, M. Kosa, and F. Hadlock. Pedagogically effective effortless algorithm visualization with a PCIL. In *Proceedings of the 39th IEEE international conference on Frontiers in Education conference*, pages 1501–1506, Piscataway, NJ, Oct 2009. IEEE Press.
- [7] T.L. Naps, S. Cooper, and twelve more authors. Evaluating the educational impact of visualization. *SIGCSE Bulletin*, 35(4):124–136, 2003.
- [8] T.L. Naps, J.R. Eagan, and L.L. Norton. JHAVÉ: An environment to actively engage students in web-based algorithm visualizations. In *SIGCSE '00: Proceedings of the thirty-first SIGCSE technical symposium on Computer science education*, pages 109–113, 2000.
- [9] T.L. Naps, G. Rössling, and nine more authors. Exploring the role of visualization and engagement in computer science education. In *ITiCSE-WGR '02: Working Group Reports from ITiCSE on Innovation and Technology in Computer Science Education*, pages 131–152, New York, NY, USA, 2002.
- [10] T.L. Naps and B. Swander. An object-oriented approach to algorithm visualization: easy, extensible, and dynamic. In *SIGCSE '94: Proceedings of the twenty-fifth SIGCSE symposium on Computer science education*, pages 46–50, 1994.
- [11] Open archives initiative - protocol for metadata harvesting - v2.0. <http://www.openarchives.org/OAI/openarchivesprotocol.html>, June 2002.
- [12] OpenAlgoViz SourceForge Site. <https://sourceforge.net/projects/openalgotv/>, 2009.
- [13] J. Preece, B. Nonnecke, and D. Andrews. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, 20(2):201 – 223, 2004.
- [14] S. Rodger and T. Finley. *JFLAP - An Interactive Formal Languages and Automata Package*. Jones and Bartlett, 2006.
- [15] G. Rössling. Translator: A package for internationalization for Java-based applications and GUIs. In *Proceedings of the 12th ACM SIGCSE/SIGCUE International Conference on Innovation and Technology in Computer Science Education (ITiCSE 2006)*, page 312, Bologna, Italy, 2006. ACM Press.
- [16] C.A. Shaffer, M.L. Cooper, A.J.D. Alon, M. Akbar, M. Stewart, S.Ponce, and S.H. Edwards. Algorithm visualization: the state of the field. *ACM Transactions on Computing Education*, 10:1–22, August 2010.
- [17] C.A. Shaffer, T.L. Naps, S.H. Rodger, and S.H. Edwards. Building an online educational community for algorithm visualization. In *SIGCSE '10: Proceedings of the 41st ACM technical symposium on Computer science education*, pages 475–476, New York, NY, USA, 2010. ACM.
- [18] J.M. Spool. The magic behind amazon’s 2.7 billion dollar question. <http://www.uie.com/articles/magicbehindamazon>, March 2009.
- [19] S.L. Toral, M.R. Martinez-Torres, F. Barrero, and F. Cortes. An empirical study of the driving forces behind online communities. *Internet Research*, 19(4):378–392, 2009.
- [20] Virginia Tech. Data Structures and Algorithm Visualization Wiki. <http://algotv.cs.vt.edu>, 2006.